

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:

The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

NEAR-PERFECT TOKEN DISTRIBUTION

by

A. Z. Broder
DEC Systems Research Center
Palo Alto, CA USA

and

A. M. Frieze
Department of Mathematics, Carnegie Mellon University
Pittsburgh, PA USA

and

E. Shamir
Department of Computer Science, The Hebrew University,
Jerusalem, Israel

and

E. Upfal
IBM Almaden Research Center, San Jose, CA, USA
University of Cambridge, Cambridge, U.K.

and

The Weizmann Institute of Science
Rehovot, Israel

Research Report No. 92-143

January, 1992

510.6
C28R
92-143

Near-perfect token distribution

A. Z. Broder * A. M. Frieze[†] E. Shamir[‡] E. Upfal[§]

January 23, 1992

Abstract

Suppose that n tokens are arbitrarily placed on the n nodes of a graph. At each parallel step one token may be moved from each node to an adjacent node. An algorithm for the *near-perfect token distribution* problem redistributes the tokens in a minimum number of steps, so that, at the end, no more than $O(1)$ tokens reside at each node. (In perfect distribution, at the end, exactly one token resides at each node.)

In this paper we present a simple algorithm that works for all *extrovert* graphs, a new property which we define and study. In terms of connectivity requirements, extrovert graphs are in-between expanders and compressors. Our results lead to an optimal solution for the near-perfect token distribution problem on almost all cubic graphs. The new solution is conceptually simpler than previous algorithms, and applies to graphs of minimum possible degree.

*DEC Systems Research Center, 130 Lytton Ave, Palo Alto, CA.

[†]Department of Mathematics, Carnegie-Mellon University. A portion of this work was done while the author was visiting DEC SRC. Supported in part by NSF Grant CCR 0089112.

[‡]Department of Computer Science, The Hebrew University. A portion of this work was done while the author was visiting DEC SRC.

[§]IBM Almaden Research Center, San Jose, CA, and The Weizmann Institute of Science, Rehovot, Israel.

1 Introduction

Suppose that n tokens are arbitrarily placed on the n nodes of a graph, with no more than M tokens at each site. At each parallel step one token may be moved from each node to an adjacent node. An algorithm for the *near-perfect token distribution* problem redistributes the tokens in a minimum number of steps, so that, at the end, no more than $O(1)$ tokens reside at each node. (In perfect distribution, at the end, exactly one token resides at each node.)

Algorithms for the near-perfect token distribution have direct applications to fundamental problems in parallel and distributed systems. We give two examples here: First, consider each node as a processor, each edge as a communication line, and each token as a message. An efficient solution to the near-perfect token distribution problem can be combined with an efficient permutation routing algorithm to give an efficient routing algorithm for an arbitrary communication request. Second, consider each token as an independent process awaiting execution and each node as a processor. Then an efficient algorithm for the near-perfect token distribution problem can be used for workload distribution among the processors.

These, and other applications, require a solution in a fully distributed model, where all information processing are done by processors that reside at the nodes of the graph. When the computation is initiated, each process has no information on the status (e.g. number of tokens) of other processors in the system. Information is transmitted via messages along the edges of the graph. In each step a processor can either send one token through an edge, or one $O(\log n)$ -bit message. Furthermore, most applications require an efficient solution on graphs that are as sparse as possible.

Peleg and Upfal [PU89] have shown that any solution to the near-perfect token distribution problem requires $\Omega(M + \log n)$ parallel steps, and this runtime can be achieved only on expander graphs. They gave an optimal $O(M + \log n)$ algorithm for *sqrt* (α, β) -*expander* graphs. These are n node, d -regular expander graphs with the following additional strong properties: there exist constants $\alpha > 1$ and $\beta > 0$ such that any set of vertices X , with $|X| \leq \beta n$, has at least $\alpha |X|$ neighbors outside X and there exists an i with $0 < i \leq d$ such that $\alpha > (d + i^2)/(i + 1)$. These two conditions imply that the expansion is at least \sqrt{d} .

In this paper we present an optimal solution for almost all cubic graphs, which is the best one can hope for from the point of view of the minimum degree required.

Our solution is based on identifying a novel property, called *extrovertness* which replaces the *sqrt*-expander property of [PU89]. Extrovertness is stronger than the expansion property but weaker than the compression property of [Pi85] (i.e every extrovert graph is an expander, and every compressor graph is an extrovert graph). In Section 2 we present an asymptotically optimal solution for any extrovert graph, which immediately leads to a solution for any graph G such that for some bounded constant c , the union $\cup_{i \leq c} G^i$ is extrovert. The main result of this paper, which is

perhaps of independent significance, is given in Section 3: we show that for almost all cubic graphs G , the graph $G \cup G^2$ (that is, the graph of all paths of length at most two in G) is extrovert. (In fact, it can be even shown that for almost all cubic graphs the graph G^2 is extrovert, but the proof is rather intricate and the constants are much worse.)

2 The New Algorithm

Our new approach to the token distribution problem is based on the *extrovertness* property:

Definition 1 A vertex $v \in S \subseteq V$ is called *extrovert in S* if v has strictly more neighbors in $V \setminus S$ than it has in S .

Definition 2 A set $S \subseteq V$ is called *b -extrovert* if at least $b|S|$ of its vertices are extrovert.

Definition 3 A graph G is called *(a, b) -extrovert* if any set $S \subseteq V$ of size at most $a|V|$ is *b -extrovert*.

Our algorithm gives an optimal solution for any n -node (a, b) -extrovert graph $G = (V, E)$.

Let d be the maximum degree in G and let $\gamma = (2d + 1)/(2d + 2)$. The algorithm works in $\log_{1/\gamma} M$ phases. (Recall that M is the initial maximum number of tokens in any node.)

Let $t_j(v)$ be the number of tokens at node v at the start of phase j , and let $M_j = \max_v \{t_j(v)\}$. Let $U_j = \{v : t_j(v) \geq M_j/2\}$ denote the set of “heavily loaded” nodes.

Each phase reduces the maximum number of tokens in any node by a factor of γ , until no vertex has more than $2/a = O(1)$ tokens. Note that in any phase $|U_j| \leq a|V|$ because $M_j > 2/a$ and the total number of tokens is $|V|$.

At the beginning of phase j the algorithm constructs a dag (directed acyclic graph) $D_j = (V_j, E_j)$ with the following properties:

1. $U_j \subseteq V_j \subseteq V$. (All heavily loaded vertices are in the dag.)
2. Each edge in D is the directed version of some edge in G , and each edge in G has at most one directed version in D .
3. For every $v \in U_j$, $\text{outdegree}(v) > \text{indegree}(v)$.

After the dag is constructed, phase j consists of $M_j/(2d + 2)$ iterations; in each iteration each node in V_j sends one token (if it has one) through each of its outgoing edges in D_j .

Lemma 1 $M_{j+1} \leq \gamma M_j$.

Proof: Let v be an arbitrary vertex in G . If v is not in D_j , and therefore not in U_j , then it does not receive (or send) any tokens in phase j , and thus it has less than $M_j/2 < \gamma M_j$ tokens at the end of this phase.

If v is in the dag but not in U_j , it starts the phase with less than $M_j/2$ tokens, and it may receive in this phase no more than $dM_j/(2d+2)$ tokens. Thus at the end of the phase it has no more than $(2d+1)/(2d+2)M_j = \gamma M_j$ tokens.

Finally, if $v \in U_j$, it sends more tokens than it receives in each iteration. Thus, it loses at least $M_j/(2d+2)$ tokens in phase j and has less than γM_j tokens at the end of the phase. \square

For a set of vertices $S \subseteq V$ let $\Phi(S) = \{v \in S : v \text{ is extrovert}\}$. Given an (a, b) -extrovert graph we construct the dag D_j , “bottom-up,” using the following procedure.

1. Set $W \leftarrow U_j$.
2. Repeat until $W = \emptyset$:
 - (a) Connect all vertices in $\Phi(W)$ to all of their G neighbors outside W .
 - (b) $W := W \setminus \Phi(W)$.

Lemma 2 *The above procedure constructs a dag with the required properties, and terminates in $O(\log N)$ iterations.*

Proof: Clearly when the procedure terminates, all of U_j is in the dag, and each edge in G was used no more than once. When a vertex $v \in U_j$ joins the dag it has more neighbors in the dag than in the remaining W , thus $\text{outdegree}(v) > \text{indegree}(v)$ in the final dag. Since the graph G is (a, b) -extrovert, and $|W| \leq |U_j| \leq a|V|$, in each iteration $|\Phi(W)| \geq b|W|$, and the procedure terminates within $O(\log n)$ iterations. \square

Corollary 3 *If G is an (a, b) -extrovert graph, then the algorithm solves the near-perfect token distribution problem on G in $O(\log M)$ phases. When phase $j-1$ ends, no vertex has more than $M_j = \gamma^{j-1}M$ tokens and phase j takes $O(\gamma^{j-1}M + \log n)$ parallel steps.*

If we execute the $O(\log M)$ phases sequentially the total run time is $O(M + \log n \log M)$. Peleg and Upfal [PU89] have shown that all the dags for the last $O(\log \log n)$ phases can be constructed in parallel. Since $\log n$ becomes dominant only when $M_j \leq \log n$, this suffices to reduce the total run-time to $O(M + \log n)$. The same idea can be applied to our algorithm. We briefly sketch the construction here, the reader is referred to [PU89] for a more detailed description.

The difficulty in constructing the dags in parallel is that in order to construct the dag D_j one needs to know the set U_j , and the set U_j is defined in turn, by the dag D_{j-1} . This difficulty is overcome by replacing the set U_j in the construction, by an overestimate set \tilde{U}_j that contains U_j and is easier to compute. We also compute, for each node, an overestimate of the number of tokens it will have at the end of each phase.

The overestimate sets are constructed *sequentially* as follows: Assume that we have already constructed the set \tilde{U}_j , and that each node has a bound on the number of tokens it will have at the end of phase $j - 1$. We run the first $c \log \log n$ phases of the construction of D_j , assuming that $U_j = \tilde{U}_j$. Let \tilde{W}_j denote the set W at the end of these $O(\log \log n)$ phases. We fix the number of phase we run so that $|\tilde{W}_j| = O(n/(\log n)^2)$.

Now we compute for each node, a bound on the number of tokens it will have at the end of phase j . If a node v has a neighbor in \tilde{W}_j , it assumes that this neighbor will be connected to it in D_j , and will send to it $\gamma^{j-1}M$ tokens. Besides this, if v is in the partially computed dag, it can compute in $O(\log \log n)$ time, a bound on how many tokens it will receive in phase j . Thus, every node which potentially is in the dag D_j , has an estimate on the number of tokens it will receive in phase j . Vertex v is in \tilde{U}_{j+1} if according to this estimate it has more than $\gamma^j M/2$ tokens, at the end of phase j .

These estimates add no more than $d|\tilde{W}_j|M_j$ “dummy” tokens in each phase. We can fix the parameters so that no more than n dummy tokens are created in the whole process, thus, this procedure can be executed until $M_j \leq 4/b$. Once the sets \tilde{U}_j are constructed, the $O(\log \log n)$ dags can be constructed in parallel in $O(\log n)$ parallel steps, with messages of size $O(\log \log n)$ bits. Thus, we prove:

Theorem 4 *The near-perfect token distribution problem can be solved on any (a, b) -extrovert graph G in $O(M + \log n)$ time.*

Let $H(G, c) = \cup_{j \leq c} G^i$, for some constant c (i.e., the graph $H(G, c)$ has an edge for any path of length c or less in G). Clearly, a parallel step on the graph $H(G, c)$ can be simulated on the original graph G in constant time. Thus, our result applies to any graph G such that $H(G, c)$ is extrovert for some constant c .

The main result of this paper, which we prove in the next section, is that if G is a random cubic graph, then with high probability the graph $H = G \cup G^2$ is extrovert. Thus we prove:

Theorem 5 *There is an $O(M + \log n)$ solution for the near-perfect token distribution problem on almost all cubic graphs.*

3 Extrovert graphs via random cubic graphs

We first state three properties that almost all random cubic graphs possess; then we show that if a cubic graph G has these properties, then the graph $H = G \cup G^2$ is extrovert.

In the definition of these properties $G = (V, E)$ is a cubic graph on n vertices, S is a subset of V , and $i(S)$ is the number of edges that are contained in S . All constants were chosen for convenience and no attempt was made to optimize them.

Property 1 *If $S \subseteq V$ and $|S| \leq (\ln n)/4$ then $i(S) \leq |S|$.*

Property 2 *For any fixed $\epsilon > 0$ there exists a function $a(\epsilon) > 0$ such that for $S \subseteq V$, if $|S| \leq a(\epsilon)n$ then $i(S) \leq (1 + \epsilon)|S|$.*

Property 3 *There are fewer than $\ln \ln n$ vertices on cycles of length less than $\ln \ln \ln n$*

We will now prove that almost all cubic graphs have these three properties. We will use the model of Bender and Canfield [BC78] and Bollobás [Bo85]. Specifically, we will adopt the notation of [Bo85].

Let $[k]$ stand for the set $\{1, \dots, k\}$. Let $W = [3n]$ and let $W_i = \{3(i-1) + 1, 3(i-1) + 2, 3i\}$, for $i = 1, 2, \dots, n$ be a partition of W into n sets of size 3, called *blocks*. For $w \in W$ we define $\psi(w) = \lceil w/3 \rceil$ so that $w \in W_{\psi(w)}$. The elements of W are called *points* to distinguish them from the vertices of the graph.

A configuration is a partition of W into $m = 3n/2$ pairs. The set of configurations is denoted Φ . For $F \in \Phi$ let $\mu(F)$ be the multigraph with vertex set $[n]$ and m edges $\{\{\psi(x), \psi(y)\} : \{x, y\} \in F\}$.

We consider Φ as a probability space in which each $F \in \Phi$ is equally likely. In this setting, for large n , the probability that $\mu(F)$ is a simple graph is at least $1/10$ and conditional on this, each cubic graph with vertex set $[n]$ is equally likely. We need only show therefore that almost all $\mu(F)$ have the above three properties.

Proof of Property 1:

$$\begin{aligned} \Pr(\text{Prop. 1 fails}) &\leq \sum_{k=4}^{(\ln n)/4} \mathbf{E}(\# \text{ of } k\text{-sets with at least } k+1 \text{ edges}) \\ &\leq \sum_{k=4}^{(\ln n)/4} \binom{n}{k} \binom{3k}{k+1} \left(\frac{k}{n}\right)^{k+1} \leq \sum_{k=4}^{(\ln n)/4} \left(\frac{ne}{k}\right)^k 8^k \left(\frac{k}{n}\right)^{k+1} \\ &= \frac{1}{n} \sum_{k=4}^{(\ln n)/4} k(8e)^k = o(1) \end{aligned}$$

Explanation of second inequality: $\binom{n}{k}$ counts sets of k blocks. Focus now on the blocks W_1, W_2, \dots, W_k . If vertex set $\{1, 2, \dots, k\}$ contains more than k edges then there are $k + 1$ pairs contained in $X = W_1 \cup W_2 \cup \dots \cup W_k$. We can choose the lowest numbered point of each pair in at most $\binom{3k}{k+1}$ ways. Having fixed these as a_1, a_2, \dots, a_k say, the probability that they are each paired in X is at most $(k/n)^{k+1}$, for given the pairings of a_1, a_2, \dots, a_i , the probability that a_{i+1} is paired in X is at most k/n . \square

Proof of Property 2: For definiteness let

$$a(\epsilon) = \left(\frac{(1+\epsilon)^{1+\epsilon}}{2e^{2+\epsilon}3^{1+\epsilon}} \right)^{1/\epsilon}.$$

Then

$$\begin{aligned} \Pr(\text{Prop. 2 fails}) &\leq \sum_{k=4}^{a(\epsilon)n} \binom{n}{k} \binom{3k}{(1+\epsilon)k} \left(\frac{k}{n} \right)^{(1+\epsilon)k} \\ &\leq \sum_{k=4}^{a(\epsilon)n} \left(\frac{ne}{k} \left(\frac{3ek}{(1+\epsilon)n} \right)^{1+\epsilon} \right)^k \\ &= \sum_{k=4}^{a(\epsilon)n} \left(\frac{e^{2+\epsilon}3^{1+\epsilon}}{(1+\epsilon)^{1+\epsilon}} \left(\frac{k}{n} \right)^\epsilon \right)^k = o(1). \end{aligned}$$

(Split the range at, say, $k = \sqrt{n}$. If $k \leq \sqrt{n}$ then $k/n \leq 1/\sqrt{n}$ so the sum up to $k = \sqrt{n}$ is $O(1/n^{2\epsilon})$; for $k \geq \sqrt{n}$ the summand is at most 2^{-k} .) \square

Proof of Property 3: Say that a cycle is small if it has at most $\ell = \ln \ln \ln n$ vertices. Then

$$\begin{aligned} \mathbf{E}(\# \text{ of vertices on small cycles}) &\leq \sum_{k=3}^{\ell} k \binom{n}{k} \frac{1}{2} (k-1)! 6^k \left(\frac{1}{3n-2k} \right)^k \\ &\leq \sum_{k=3}^{\ell} 2^k = o(\log \log n). \end{aligned}$$

and so we can use the Markov inequality. (Explanation of first inequality: having chosen k vertices and a cyclic ordering in $\binom{n}{k} \frac{1}{2} (k-1)!$ ways, there are 3^k choices for the points in pairs not in edges of the cycle and 2^{k-1} ways of pairing up the points in the blocks to make the edges of the cycle. The probability that these k pairs exist is at most $(3n-2k)^{-k}$.) \square

Theorem 6 *If G is a cubic graph that has properties 1, 2, and 3, then the graph $H = G \cup G^2$ is extrovert.*

Proof: For a set $S \subseteq V$ let

$$\Phi(S) = \{v \in S : v \text{ is extrovert in } S \text{ with respect to } H\}$$

and let $g(S) = |\Phi(S)|$.

From now until the end of the proof all adjacency and neighborhood relations, when not specified otherwise, are with respect to G .

Let ϵ be a sufficiently small constant and let S be a set of vertices with $|S| \leq \frac{1}{4}a(\epsilon)n$. Let $T = \Gamma(S)$ be the set of neighbors of S in G . (Thus $|S \cup T| \leq a(\epsilon)n$.) Let $i = i(S)$, the number of edges in S . Now partition S and T as follows: Let

$$\begin{aligned} S_j &= \{v \in S : v \text{ has } j \text{ neighbors in } S\}, & j &= 0, 1, 2, 3. \\ T_j &= \{v \notin S : v \text{ has } j \text{ neighbors in } S\}, & j &= 1, 2, 3. \end{aligned}$$

Finally let $s = |S|$, $s_j = |S_j|$, and $t_j = |T_j|$.

We first prove that S is extrovert in H under two assumptions:

1. No vertex of S lies on a cycle of length ≤ 4 in G .
2. $|S| \geq s^* > 0$ where s^* is an absolute constant.

These assumptions are removed in the last part of the proof.

Under Assumption 1, each vertex in V has 9 distinct neighbors in H . For v to be extrovert with respect to S in H , it must have at least 5 neighbors in H outside S . Having two G -neighbors in T_1 is sufficient, but if v has no neighbor in S_3 , then one G -neighbor in T_1 is sufficient.

There are two cases in the proof. In the first case, $i \geq (1 - \alpha)s$ for a suitable chosen $\alpha > 0$. In this case $T_2 \cup T_3$ is quite small, and the proof works readily; In the second case, $i < (1 - \alpha)s$, and the graph induced by S breaks into $k \geq \alpha s$ components out of which $k' \geq (1 - 2\alpha)k$ are trees of size at most α^{-2} . Here we estimate the number of extrovert vertices in terms of αk .

We start by deriving a series of relations among the s_j 's and the t_j 's.

From Property 2 applied to $S \cup T_2 \cup T_3$ we deduce that

$$i + 2t_2 + 3t_3 \leq (1 + \epsilon)(s + t_2 + t_3),$$

or

$$(1 - \epsilon)t_2 + (2 - \epsilon)t_3 \leq (1 + \epsilon)s - i. \tag{1}$$

Let

$$S_2^g = \{v \in S_2 : v \text{ has no neighbor in } S_3\}$$

and let $s_2^g = |S_2^g|$. Observe that if $v \in S_2^g$ is adjacent to $w \in T_1$ then $v \in \Phi(S)$. Clearly

$$s_2^g \geq s_2 - 3s_3 \tag{2}$$

Claim 7

$$s_1 \geq s_3 - 2\epsilon s$$

Proof:

Consider the graph G_S induced by S . Let C be a component of G_S with at least one edge. Let C have σ_i vertices of degree i , for $i = 1, 2, 3$, and let $\sigma = |C|$. Due to Property 2, by removing at most $\epsilon\sigma + 1$ edges from C we can obtain a spanning tree C' . Observe that in C'

$$\sigma'_1 = \sigma'_3 + 2.$$

Now add back the removed edges. Each addition reduces the difference between the number of degree 3 vertices and degree 1 vertices by at most two. Hence

$$\sigma_1 - \sigma_3 \geq \sigma'_1 - \sigma'_3 - 2(\epsilon\sigma + 1) = -2\epsilon\sigma.$$

Now add up over all the components of S . \square

Now from equations (2) and Claim 7 we obtain that

$$\begin{aligned} s &= s_0 + s_1 + s_2 + s_3 \\ &\leq s_0 + s_1 + s_2^g + 4s_3 \\ &\leq s_0 + s_1 + s_2^g + 4s_1 + 8\epsilon s \end{aligned}$$

Hence

$$s_0 + 5s_1 + s_2^g \geq (1 - 8\epsilon)s. \quad (3)$$

Now choose α small but large compared with ϵ . There are two cases.

Case 1: $i \geq (1 - \alpha)s$.

We deduce from (1) that

$$2t_2 + 3t_3 \leq \frac{2}{1 - \epsilon}(\alpha + \epsilon)s \leq 3\alpha s. \quad (4)$$

The number of edges coming from $T_2 \cup T_3$ into S is $2t_2 + 3t_3$. Delete from S all the vertices hit by such edges. A remaining vertex v has out-of- S neighbors only in T_1 , thus v has at least two such neighbors if $v \in S_0 \cup S_1$ and at least one such neighbor if $v \in S_2^g$. Hence

$$\begin{aligned} g(S) &\geq s_0 + s_1 + s_2^g - (2t_2 + 3t_3) \\ &\geq \frac{1}{5}(s_0 + s_1 + s_2^g) - 3\alpha s \\ &\geq \left(\frac{1 - 8\epsilon}{5} - 3\alpha\right)s \end{aligned} \quad (5)$$

Case 2: $i < (1 - \alpha)s$.

In this case the number of components k of G_S satisfies

$$k \geq \alpha s. \quad (6)$$

The number of components of G_S of size greater than $s/(\alpha k)$ is at most αk . Hence there are at least $(1 - \alpha)k$ components of size $< s/(\alpha k) \leq \alpha^{-2}$.

Claim 8 *If $i < (1 - \alpha)s$ then the number of components of S of size $\leq \alpha^{-2}$ that are trees is at least $(1 - 2\alpha)k$.*

Proof: We can assume without loss of generality that $G_{S \cup T}$ is connected.

If $s + t \geq A \ln n / \ln \ln n$ then we use Property 3 to assert that at most $o(k)$ of the small components are not trees.

If $s + t < A \ln n / \ln \ln n$ then at most one of the small components is not a tree; otherwise two non-tree components can be connected via T and contradict Property 1. By Assumption 2

$$s \geq s^* > \alpha^{-2} \quad (7)$$

so $\alpha k \geq \alpha^2 s \geq 1$.

Hence in both cases the number of small components that are not trees is less than αk . \square

Suppose now that the k components have sizes c_1, c_2, \dots, c_k , and that these components span i_1, i_2, \dots, i_k edges. We re-write equation (1) as

$$\begin{aligned} (1 - \epsilon)t_2 + (2 - \epsilon)t_3 &\leq (1 + \epsilon) \sum_{1 \leq t \leq k} c_t - \sum_{1 \leq t \leq k} i_t \\ &= \sum_{1 \leq t \leq k} (c_t - i_t) + \epsilon s \leq k + \epsilon s \end{aligned} \quad (8)$$

Now we divide the components of s of size $\leq \alpha^{-2}$ that are trees into three classes:

- k_1 trees that are isolated vertices.
- k_2 trees that are paths of length ≥ 1 .
- k_3 trees that have at least one degree 3 vertex.

We now classify the leaves of these trees. Within these trees there are

- l_1 isolated vertices adjacent to T_1 , and thus $k_1 - l_1$ other vertices triply adjacent to $T_2 \cup T_3$.

- l_2 end-of-path leaves adjacent to T_1 , and thus $2k_2 - l_2$ other vertices doubly adjacent to $T_2 \cup T_3$.
- l_3 other leaves doubly adjacent to T_1 , and thus at least $3k_3 - l_3$ other vertices adjacent to $T_2 \cup T_3$.

Note that $l_1 + l_2 + l_3$ counts extrovert vertices with respect to H . Counting edges out of $T_2 \cup T_3$ we get the inequality

$$2t_2 + 3t_3 \geq 3(k_1 - l_1) + 2(2k_2 - l_2) + (3k_3 - l_3)$$

or

$$3l_1 + 2l_2 + l_3 \geq (3k_1 + 4k_2 + 3k_3) - (2t_2 + 3t_3)$$

and hence

$$\frac{1}{3}(l_1 + l_2 + l_3) \geq (k_1 + k_2 + k_3) - \frac{2}{3}(t_2 + 2t_3).$$

Now using Claim 8 and equation (8) and assuming $\epsilon < 1/3$, we obtain that

$$\begin{aligned} \frac{1}{3}(l_1 + l_2 + l_3) &\geq (1 - 2\alpha)k - \frac{1}{1 - \epsilon} \frac{2}{3}(k + \epsilon s) \\ &\geq \left(\frac{1}{3} - 2\alpha\right)k - \epsilon s \geq \left(\left(\frac{1}{3} - 2\alpha\right)\alpha - \epsilon\right)s. \end{aligned}$$

Hence

$$g(S) \geq ((1 - 6\alpha)\alpha - 3\epsilon)s/3 \tag{9}$$

Equations (5) and (9) conclude the proof under the two initial assumptions, since they show that for any set $S \subseteq V$, if $|S| \leq a(\epsilon)n$, then $g(S)/|S|$ is greater than a certain constant. We now proceed to remove the assumptions that we made.

Let b be the lower bound on $g(S)/|S|$ derived above. Suppose first that $|S| \geq (\ln n)/4$. Remove the $k \leq \ln \ln n$ vertices in S that lie on cycles of length $\leq \ln \ln \ln n$, thus obtaining S' . All extrovert vertices in S' are extrovert in S with the possible exceptions of vertices at distance 1 or 2 from the removed vertices. (Because the removed vertices are on cycles, there are only $3k$ neighbors.) Therefore

$$g(S) \geq g(S') - 3k \geq (b - o(1))|S|.$$

Suppose now that

$$\left(8 + \frac{24}{b}\right) \leq |S| \leq \frac{1}{4} \ln n$$

By Property 1, there is at most one cycle in $S \cup \Gamma(S) \cup \Gamma^2(S)$. If this cycle has length less than 4 and contains vertices of S , remove them to obtain S' . Then, as before

$$g(S) \geq g(S') - 12 \geq b(s - 4) - 12 \geq \frac{bs}{2}.$$

It remains to consider the case $2 \leq |S| \leq s^*$, where $s^* = \max\{8 + 24/b, \alpha^{-2}\}$, to accommodate also equation (7). We shall argue that S contains at least one extrovert vertex.

Consider the graph K that consists of all edges incident to S in G . (In other words, K is the graph induced by $S \cup T$ except for the edges whose both endpoints are in T .) Without loss of generality we can assume that K is connected.

Clearly K has the following properties:

- Any $v \in S$ has degree 3 in K .
- Any leaf in K belongs to T_1 .
- Any neighbor of a leaf in K belongs to S .
- There is at most one cycle in K .

If K is a tree then there must be a $v \in S$ that is the neighbor of two leaves in K , and therefore is extrovert in S . Similarly if K contains a cycle, then if there is any $v \in S$ that does not belong to the cycle, then there must be a $v' \in S$ that is the neighbor of two leaves in K . (Consider the tree that “hangs” at v away from the cycle.) It remains to consider the case when all the vertices in S belong to a cycle in K . But then all vertices in S are extrovert. \square

References

- [BC78] E. A. Bender and E. R. Canfield, The asymptotic number of labeled graphs with given degree sequences. *Journal of Combinatorial Theory B*, 24:296–307, 1978.
- [Bo85] B. Bollobás, Random graphs, Academic Press, 1985.
- [Pi85] N. Pippenger, On networks of noisy gates. *26th Annual Symposium on Foundations of Computer Science*, pp. 30–38, Portland, Oregon, October 1985.
- [PU89] D. Peleg, and E. Upfal. The token redistribution problem. *SIAM Journal of Computing*, 18:229–243, 1989.
- [PU87] D. Peleg, and E. Upfal. The Generalized Packet Routing Problem. *Theoretical Computer Science*, 53:281–293, 1987.

FEB 13 2004

Carnegie Mellon University Libraries



3 8482 01373 5671